

FACILITATING MAINTENANCE OF INDEXES DURING A REORGANIZATION OF DATA IN A DATABASE

Wayne E. Fisher

ABSTRACT OF THE DISCLOSURE

5 An index can be maintained during a reorganization of data in a database by retaining each root segment, or the prefix component of each root segment, in its storage location during the reorganization. Correction of an index after a reorganization of data in a database can be facilitated by assigning a unique token to each target segment, or the prefix component of each target segment, and each
10 corresponding index entry having an address to a target segment, prior to a reorganization of data. The unique token for a given target segment, or the prefix component for a given target segment, and the unique token for a corresponding index entry are the same. After a reorganization of data in the database, the unique token of a first index entry is read. Then, the unique token of each target segment, or the
15 prefix component of each target segment, is read until a match is found between the unique token of a matching target segment, or the prefix component of a matching target segment, and the unique token of the first index entry. After a match is found, the address of the first index entry is replaced with the address of the matching target segment, or prefix component of the matching target segment. If the first index entry
20 is associated with a particular database record, the search for a token which matches the first index entry's token can be focused on the segments for that database record. All the segments for a record can be stored within a single block of storage locations. Before the address of an index entry is corrected, it can be determined if the address is valid. The address of the index entry will then be corrected only if it is invalid.